arXiv:1402.2793v2 [cs.MA] 21 May 2014

NOTICE: this is the author's version of a work that was accepted for publication by Elsevier. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in Future Generation Computer Systems, Volume 37, July 2014, Pages 390–400, http://dx.doi.org/10.1016/j.future.2014.02.002 .

Computing Agents for Decision Support Systems

D. Krzywicki[a], Ł. Faber[a], A. Byrski[a,*], M. Kisiel-Dorohinicki[a]

[a] *AGH University of Science and Technology, Al. Mickiewicza 30, 30-059 Krakow, Poland*

Abstract

In decision support systems, it is essential to get a candidate solution fast, even if it means resorting to an approximation. This constraint introduces a scalability requirement with regard to the kind of heuristics which can be used in such systems. As execution time is bounded, these algorithms need to give better results and scale up with additional computing resources instead of additional time. In this paper, we show how multi-agent systems can fulfil these requirements. We recall as an example the concept of Evolutionary Multi-Agent Systems, which combine evolutionary and agent computing paradigms. We describe several possible implementations and present experimental results demonstrating how additional resources improve the efficacy of such systems.

Keywords: Decision Support Systems, Multi-agent Systems, Scalability, Performance

1. Introduction

The need to gather and analyse vast amounts of information from numerous sources has grown in importance. Such data is often a basis for simulations and computations that support decision making. It may be needed to run many computing tasks, in order either to test different parameters in a model or to verify a statistical hypothesis. An exhaustive search for optimal solutions to a decision making problem is usually time-consuming and thus not acceptable in real-time conditions. Instead, metaheuristics may quickly provide good-enough options to be further considered in the decision making process [1].

Examples of use cases with real-time constrains, where a quick approximated solution may be better than an outdated optimal one, may include:

- Portfolio optimisation — a decision support system can apply different models to the available market data and allow the user to quickly react to arising trends [2].

*Corresponding author

Email addresses: daniel.krzywicki@agh.edu.pl (D. Krzywicki), faber@agh.edu.pl (Ł. Faber), olekb@agh.edu.pl (A. Byrski), doroh@agh.edu.pl (M. Kisiel-Dorohinicki)

- Crisis management — in crisis situations, such as fire outbreaks, flooding or earthquakes, intensive simulations are required in order to suggest possible evacuation routes or to assign rescue units to tasks. Geographical information usually needs to be considered, yielding optimisation problems similar to transportation-related ones [3].
- Production planning — decision support systems can help in scheduling work, rescheduling production plans in case of hardware failures, implementing just in time strategies or balancing conflicting goals (e.g. high system throughput vs low machinery usage) [4].

Metaheuristics may still require a significant computational power if the acceptable solution is to be found in a reasonable time. For this purpose, large-scale infrastructure is usually used, such as clusters, grids or clouds. To fully benefit from this computational power, it is required to appropriately plan their development and deployment, along with adequate tools and careful testing.

Because of its intrinsic decentralisation [5], the agent approach is well suited to design scalable distributed models and has been applied in various decision support systems. This approach may be summarised as the introduction of artificial intelligence techniques into the system, transforming it from a passive tool into an active collaborator in decision making. A number of such case-oriented systems have been proposed and verified in practice [6, 7].

Well-known general-purpose agent-based development tools (such as JADE [8], RePast [9] or Madkit [10]) may not be the best choice to implement such computational intensive simulations, when throughput and scalability are more important than code migrations or FIPA-compliant communication. Therefore, over the last 10 years, we have been involved in the development of several alternative platforms dedicated to large scale agent-based simulations and computations [11, 12, 13].

In this work, we discuss the implementation aspects of using computing agents in large-scale environments, with a focus on performance. We compare different approaches to agent execution and parallelism, based on the metaheuristic called evolutionary multi-agent systems (EMAS), which is a hybrid of agent-oriented and evolutionary-based computing [14]. We introduce the concept of *meeting arenas*, which allow to design more efficient and scalable multi-agent systems. Nevertheless, we show that explicit parallelism, when each agent is mapped onto a thread, can be much less effective than a simple but optimised sequential implementation. Finally, we show that such agent-based metaheuristics can be easily scaled with additional computational resources.

We start the paper with a discussion on the applicability of the agent-oriented paradigm and metaheuristics in decision support systems (Section 2), along with an EMAS example. In Section 3, we introduce the most common approaches to parallelism in agent-oriented computing and follow with a review of popular agent platforms in Section 4. We describe in Section 5 how to implement an evolutionary multi-agent system using two different approaches—a synchronous and asynchronous one. Finally, we conclude the paper by comparing the performance and scalability of both approaches in Section 6.

2. Agent-Based Metaheuristics in Decision Support

Decision Support Systems (DSS) are information systems that support different business or organisational activities involving decision-making. They are especially useful in situations where quickly changing, hard to specify in advance conditions are encountered. Referring to Power's taxonomy for DSSs [15] this paper focuses on Model-driven DSSs, which help the users in the analysis of the current situation by allowing to manipulate statistical, simulational or optimisational models.

2.1. Metaheuristics for DSSs

The models used in DSSs are usually very complex and computationally hard, because the underlying problems are very difficult as well. In such cases, one often turns to solutions based on so-called heuristic methods, which provide "good-enough" solutions without caring whether they may be proved to be correct or optimal [1]. These methods trade-off precision, quality and accuracy in favour of smaller execution time and computational effort. They are necessary to deal with difficult problems, and are referred to as methods of the last resort [16].

A general definition of a heuristic algorithm, which does not specify details such as a particular problem, search space or operators, is called a *metaheuristic*. For example, a simple algorithm such as greedy search may be defined without going into more details as "an iterative, local improvements of a solution based on random sampling"[17].

A simple but adequate classification of metaheuristics (cf. [18]) distinguishes two groups of techniques. Single-solution metaheuristics work on a single solution to a problem, seeking to improve it. The examples are greedy search, tabu search or simulated annealing. Population-based metaheuristics explicitly work with a population of solutions and put them together in order to generate new solutions. The examples are evolutionary algorithms, immunological algorithms, particle swarm optimisation, ant colony optimisation, memetic algorithms and other similar techniques. They are usually inspired by nature and imitate different phenomena observed in e.g., biology, sociology, culture or physics [19].

2.2. Agent Approach

The key concept in multi-agent systems (MAS) consist in intelligent interactions, such as coordination, cooperation, or negotiation. Therefore, multi-agent systems are ideal in representing problems which can be solved using multiple methods by numerous entities with various perspectives. One of the most important features in a multi-agent system is the autonomy of the agents, as they can fulfil the tasks assigned to them according to their own strategy and the situation observed in their environment. In consequence, agents are adaptable and proactive [5].

Combining the agent-oriented approach with population-based metaheuristics seems natural but has yet been the topic of little work. The entities processed in the course of the computation can often be considered autonomous and treated as agents in a common environment. The operations involving many such entities can be defined as interactions between these agents.

This change of modelling perspective allows to perceive parts of the system on a higher abstraction level and build hybrid systems which combine techniques from different metaheuristics. New problems often arise, such as the lack of global knowledge or the need for proper synchronisation of the agents' actions. However, interesting and useful effects also often result from the cooperation of different mechanisms in one system [20].

In this paper, we focus on an example of such a hybrid approach, in which agents are subject to an evolutionary process. Such a combination yields interesting new features when compared to classical evolutionary algorithms, such as a decentralised an emergent selective pressure.

2.3. Evolutionary Multi-Agent Systems

Agents in an evolutionary multi-agent system (EMAS) represent solutions to a given optimisation problem.

Inheritance is achieved through reproduction, with the possible use of variation operators such as mutation and recombination, like in classical evolutionary algorithms. Yet agents are to be autonomous in their decisions and no global knowledge is available to them. Therefore, in contrast to classical evolutionary algorithms, selection needs to be decentralised and involve peer-to-peer interactions instead of being system-wide.

In order to do that, a solution based on the acquisition and exchange of non-renewable resources has been proposed in [21]. The quality of the solution represented by the agent is expressed by the amount of resources the agent owns. In general, these resources should pass from worse agents to better ones. This might be realised through encounters between agents, which cause better ones to end up with more resources and make them more likely to reproduce. Worse agents lose resources which increases the probability of their death. Because of such indirect dynamics of reproduction and death, agents' lifespans overlap and so do the generations. Moreover, the size of the population is dynamic and can be changed by varying the amount of available resources. A detailed study of computing with EMAS, in particular the influence of its different parameters on the computing efficiency may be found in [22].

Agents are grouped within *environments* which define the information and resources an agent has access to. Agents can interact with each other directly only within the same environment. However, they are able to move to another environment, thus exchanging information and resources all over the system [14] (see Fig. 1).

Environments are largely independent and communicate only through agent migrations. Therefore, they can be easily treated as basic units of distribution, as in the classical island model in evolutionary algorithms. In addition to

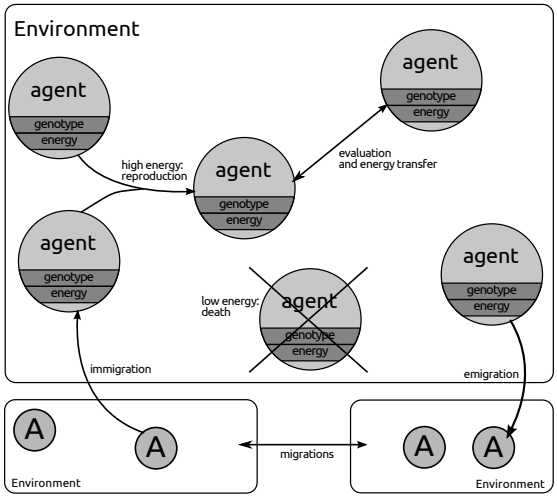

Figure 1: An example of a simple evolutionary multi-agent system (EMAS). Agents with higher fitness take energy from agents with lower fitness. High levels of energy increase the probability of reproduction and reduce chances of death. In consequence, the selection process is decentralised and selective pressure softens. Multiple agent environments can be connected through agent migrations, like in the classical island model.

improving the performance of the algorithm, it also increases the diversity of solutions in the whole population (allopatric speciation). Other metaheuristics can also be introduced, such as immunological selection [23] and niching [24].

The principle of an evolutionary multi-agent system consist in the explicit hybridisation of agent-oriented and evolutionary computing. This contrasts with usual agent-oriented approaches, which use the agent-paradigm to solve certain tasks by delegating them to particular agents and combining the outcomes of their work (see, e.g. [25]).

3. Interaction and Execution Models for Agents

In agent-oriented computing systems, agent interactions are one of the crucial aspects of their work. It is easy to predict that parallelising them can significantly increase the throughput of the system. However, this comes at the cost of increased communication and synchronisation. Therefore, an important issue is to choose the appropriate granularity of the entities in the computation.

As agents are defined as autonomous and independent beings, it seems natural to look for further concurrency within a single environment. The question is where to put the boundaries of concurrent execution, as it has consequences on both performance and ease of programming. This section discusses the most common models of execution and interaction in existing agent software.

3.1. Heavyweight Agents

In this model every agent is associated with a thread and communicates through message passing. Some agents may passively wait for incoming messages and react to them. Other agents may actively initiate interactions with other agents. It is difficult to achieve a coordinated life cycle among such agents, since the corresponding threads may be arbitrary interleaved. Therefore, some kind of synchronisation between agents still needs to be introduced, usually in terms of a specific communication protocol.

In order to interact with each other, agents need to locate other agents willing to perform the same actions. For example, in an evolutionary multi-agent system, an agent with enough resources to reproduce needs to find another one which also has enough resources. In order to do that, it could ask all other agents in the population. However, such a solution is obviously inefficient, because of the intensity and redundancy of the required communication.

A better approach, introduced in this paper, is to use a mediating entity, which we call a *meeting arena*. Every time an agent wants to perform an action, it chooses an appropriate arena to meet with other similar agents. The arena is then able to partition its members in groups of some given arity and mediate the meeting itself (see Fig. 2). Examples with pseudocode are given in Section 5.

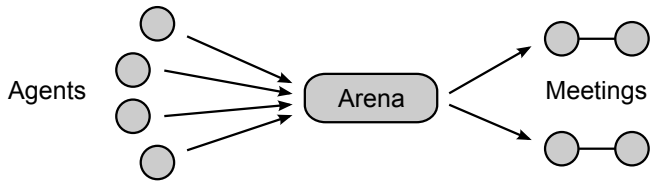

Figure 2: Meeting arenas allow to group similar agents and coordinate meetings between them.

The usage of meeting arenas should bring many benefits, not only in terms of efficiency, as the algorithm itself can be structured more clearly. Agents only need to be given a set of rules, in order to choose an arena on the basis of their state. The actual protocol of agents interactions can then be defined at the level of the appropriate arena.

Assigning a thread to each agent may feel very natural. In practice, however, the number of agents is often much higher than the number of cores, especially in simulations. Performance may then be seriously hindered by frequent context switches, although this overhead may be reduced by sharing a pool of threads among agents. However, this model still involves intensive communication and costly processor cache synchronisation. In consequence, the trade-off for such concurrency may be higher than expected.

3.2. Lightweight Agents

An opposite approach is to consider agents as parts of the model, not parts of the implementation. As such, they are simply represented as data structures and processed like in a discrete event simulation.

The execution of an individual agent has to be divided into smaller parts which can be interleaved. These parts, which we will call *actions*, could for example consist in executing a single step or querying a neighbour. Given its current state, every agent decides which action to perform next. This action is submitted for later execution to an executor service owning a pool of threads, like in the Command Object design pattern.

The difference with regard to classical discrete event simulation is that actions are generated synchronously but can be executed asynchronously. In other words, the state of an agent during the execution of the action may be different from the time when the action was created.

The performance of such a model will usually be higher than in the previous one, more consistent memory access patterns resulting in more efficient processor usage. Even though the explicit parallelism is reduced, throughput can be improved, because frequent agent interactions no longer need to be synchronised between threads.

Moreover, independent actions can still be executed in parallel by the executor service. This is consistent with the meeting arena concept described above, as actions on common subsets of agents may be grouped together and considered as a single meeting.

4. Distributed and Parallel Multi-Agent Frameworks

This section provides an overview of existing multi-agent frameworks. In our review, we focus on parallelisation and distribution capabilities, with regard to the aspects discussed in the previous section.

First, we briefly describe some selected tools which specialise in metaheuristics. They are interesting examples of improving metaheuristics with agent systems, but lack more general agent-oriented features. However, all of them share the idea of an agent being an *executor* of the algorithm and not a *participant*.

The platforms described in the consecutive sections provide more sophisticated support for agent-based systems but are not necessarily well suited to metaheuristic computations. Some properties are shared by almost all of them, such as: the choice of an object-oriented programming language (mostly Java) and a representation of agents as objects with an internalised state. Other characteristics include: models that are too heavy (e.g., JADE due to FIPA compatibility), a large and complicated code base due to the implementation of many communication, distribution and component-oriented mechanisms in the platform instead of using ready solutions (e.g., Jadex).

4.1. Metaheuristics Frameworks

The four frameworks presented below are examples of introducing multi-agent systems to metaheuristic computations. They are specialised to this purpose and lack well-defined distribution facilities. We present them as alternative models but we do not discuss their implementations in-depth.

MAGMA (Multi-AGent Architecture for Metaheuristics). MAGMA [26] is a multi-level (hierarchical) architecture of a multi-agent system. Each level of agents has different objectives and represents a different level of abstraction of the algorithm. For example, level 0 agents generate a sample solution and then level 1 agents improve it by searching the neighbourhood of that solution. There may be several agents that participate in an algorithm on each level. Composing different metaheuristics is also possible with a coordination provided by a higher level of the architecture (level 3). This way agents wrap selected functions of the algorithm and not the whole algorithm itself (as it will be the case for further platforms).

MAS-DGA (Multi-Agent System for Distributed Genetic Algorithms). The attention of the authors in [27] is focused on approaches to the question of migration. They propose the MAS-DGA framework that comes from the concept of Distributed Genetic Algorithms, where the population is divided into interacting subpopulations handled by different genetic algorithms (GA). In the case of MAS-SGA, these GA are encapsulated in agents. The authors suggest a possibility of distribution on the agent level but they do not provide descriptions of any specific examples nor implementations of this model.

AMF (Agent Metaheuristic Framework). The authors of AMF [28] extend metaheuristics with an agent-oriented approach and an organisational model based on *roles* and *interactions*. The RIO meta-model described in the paper involves the three following concepts: Role, Interaction and Organisation. Metaheuristics are organisations, and agents play specific roles in these organisations. Some of the defined roles are: the intensifier (performs a search in a search space), the diversifier (identifies new promising regions in the search space), the guide (structures the information from two previous roles), etc.

MAS4EVO (Multi-Agent System for EVolutionary Optimization). In [25] the authors propose a model and a framework (DAFO – Distributed Agent Framework for Optimization) that is a significant improvement over the previous three. The framework is built on MadKit (see Section 4.3). In this model, authors introduce three types of agents: problem solving agents which optimise functions, fabric agents which are responsible for initialising and configuring the computation, and observing agents which generate output for the end-user.

4.2. Jadex

Jadex[1] is an agent-based programming framework that exploits a novel approach to agents-components unification called "active components" [29].

The concept of Active components unifies SCA (Service Component Architecture) components with agents. This results in components that are able to use, in addition to traditional required and provided service interfaces, asynchronous messaging and that can act autonomously. It has a tremendous impact

[1] `http://www.activecomponents.org`

on behaviours of these components. They, for example, can refuse service call execution when they cannot or do not want to process the request.

Agents. Jadex offers two ways to implement agents. It is possible to use full-featured, BDI (belief-desire-intention) agents and simple, so-called *micro agents*. Micro agents are usually just annotated POJO (Plain Old Java Objects) classes. They follow three-phased execution semantics: initialisation, execution and termination. Additionally, an agent can schedule actions to be run later. As an agent is also an active component, it may receive service calls and incoming messages.

Distribution. Distribution in Jadex is provided transparently to the developer and it is implemented using a layered architecture. Services, for example, may use remote asynchronous method calls. Transparency is achieved by using proxy interfaces implementation. Internally, remote calls are implemented using asynchronous messaging between remote management system components. Messages are encoded and transmitted through some chosen *stream*. The encoding of messages is provided by *codecs* which need to (un)marshall Java objects to binary or XML format but which can also provide more sophisticated features: e.g., compression or encryption. A stream can use different communication transport: TCP, HTTP and others.

The second aspect of distribution is peer awareness and discovery. Jadex takes care of it automatically on all levels, including service (i.e. interface) binding. When there is a look-up for a required service, proxy components on a local node redirect search requests to the remote management system to perform remote look-ups of services.

Before that, remote platforms in the network need to be discovered. For this purpose, Jadex provides a few different mechanisms: e.g., broadcast discovery which sends UDP announcements about a platform on a local network or registry discovery in which there is one, central registry created for all platforms to announce themselves.

Other features of Jadex include, among other things, support for interaction with external systems using web services and a GUI-based control centre.

4.3. MadKit

MadKit[2] is a generic, customisable multi-agent platform based on a specific organisational model [10]. Agents are divided into groups and they may have particular roles in them. The centralisation of the platform around the organisational concepts is, in the view of the authors of the platform, a key element for building heterogeneous systems.

The MadKit architecture is built on the agent-group-role (AGR) model. This model is used to built organisations: an *organisation* is described using terms of interacting *groups* and *roles* and is separated from the concept of an agent.

[2] `http://www.madkit.net`

Agents. A MadKit agent is an entity that can communicate and which has several roles within one or more groups. Groups are atomic structures aggregating agents and they can overlap. Roles are tags for agent functions within groups. Agents request them on their own and they may be granted or denied them. Communication between agents is achieved using asynchronous messaging. Addressing of agents is done using their addresses or by their specific roles in one of their groups.

Architecture and Distribution. The architecture of MadKit is developed around the AGR concept. Moreover, it follows some additional design decisions, the most interesting being: micro-kernel architecture and agentified services.

The micro-kernel, responsible for basic platform management, handles only most essential functions. The rest of the needed services is handled by agents. The micro-kernel tasks are: control of groups and roles, life-cycle management of agents, local messaging. It also supports so-called "kernel hooks" which allow extension of its functionality by operations executed in the publish-subscribe model. Two types of hooks are supported: monitor and interceptor hooks. The former can be used by many agents at the same time whilst the latter can be hold on by only one agent and can be used to prevent the operation from successful execution. Additionally, it is possible to execute actions on kernel when an agent is a member of the *system* group.

The agentification of services describes a concept of turning system services (e.g., distributed message passing, migration) into agents. This makes the platform very extensible and flexible as every component can be easily replaced: communication with services is no different than with other agents.

MadKit has support for transparent distribution. Groups can span across many platform nodes. It is provided by two roles in the *system* group: communicator (routes messages to other nodes) and synchroniser (keeps information about groups memberships synchronised across all nodes).

MadKit provides also a graphical environment for visualising simulations and controlling the platform.

4.4. μ^2

μ^2 (micro-squared)[3] is a multi-agent platform centred around the concept of a μ-agent (or micro-agent): a small-size agent, that can be recursively constructed from other micro-agents with decomposition and fine-grained separation of concerns in mind [30]. μ^2 is implemented in Java and in Clojure and available under GPL 3.0 license.

Agents. Micro-agents are autonomous, persistent, reactive and proactive. They can play one or more roles which fulfil so-called applicable intents. Intent is another name for an intention or "abstract request specification". The platform provides some organisational modelling approaches. An agent can be in a

[3] `http://sourceforge.net/apps/mediawiki/micro-agents/`

group leader role. In such case, it controls many sub-agents, propagates control messages and structures the society. Due to this approach agents can construct themselves using sub-agents, and sub-agents also can be group leaders. This is how decomposition can be implemented in the platform. Other roles include: social roles that allows agents to communicate asynchronously and passive roles which support only synchronous communication. The latter role was introduced to reduce possible performance penalties resulting from asynchronous messaging.

Distribution. As it is mentioned in [30], μ^2 is a platform that can be run on Android devices. Micro-agents are encapsulated into Android service and they are integrated into the rest of the system. Communication between normal applications and agents is transparent.

4.5. JADE

JADE[4] is a mature (founded in 2000) Java framework for developing agent-based applications with a very strong relationship with FIPA specifications [8]. Its architecture is focused on a peer-to-peer communication with some centralised services. Two software components are specified: agents (autonomous, using asynchronous messaging) and services (non-autonomous, running on a single or multiple nodes).

Agents. JADE agents exist in containers (basically Java processes) which can be distributed over the network. The way messages are structured is compliant with FIPA Agent Communication Language.

Distribution. The peer-to-peer nature of JADE made it possible to create many reimplementations of nodes, e.g., for mobile environments like Android [31]. Distribution is gained by splitting a JADE container into a frontend and backend. The former runs on a mobile device and is rather lightweight, the latter usually runs on a more powerful computer.

4.6. Repast Suite Family

Repast[5] is an open-source, agent-based modelling and simulation toolkit [9]. It has many versions for various programming languages. The most interesting ones are the newest: Repast Simphony (for Java) and Repast HPC (for C++). All of them uses the "new BSD" license.

Repast Simphony is a complete rewrite of older Repast 3 with a modular architecture, extendable via plugins. Individual components (e.g., networking, logging) can be replaced easily. Plugins are layered and separate layers can be replaced with similar easiness. There is a separation between model specification, execution, data storage and visualisation.

[4] `http://jade.tilab.com/`
[5] `http://repast.sourceforge.net`

A core of the Repast Simphony consists of components responsible for simulation functions (e.g., time scheduling, space management, random number generators). ss

Agents. Agents are modelled as objects, collections of agents – as contexts, and the environment – as projections. A *context* is a set of objects and may represent an agents' population but does not describe any structure or relationships between agents. The second term – *projection* – was created to define structures of agents in contexts. They may be, for example, network or grid structures.

Distribution. Repast does not offer distribution facilities similar to other platforms. However, a user can prepare its own distributed environment using external facilities, for example, Java RMI (Remote Method Invocation), which is an object-oriented remote procedure call mechanism.

5. Implementation Aspects

Evolutionary multi-agent systems and similar agent-based computing systems need lightweight, reusable and easy-to-parallelise solutions. In particular, the implicit agent-orientation perceived at the implementation level of these platform does not seem inevitable to us. We think that agent features should be a part of the conceptual level, but do not need to be reflected in the implementation in the case of computing systems.

Considering the execution models described in section 3 and their implementation in existing software tools for multi-agent systems, we wanted to compare these two approaches to tell what granularity is best suited for agent-based computing and simulation. In order to abstract from the properties of these frameworks not relevant to the problem, we implemented two custom versions of an evolutionary multi-agent system.

In the first version agents are asynchronous and can be mapped to separate threads or share a thread pool. The second version is synchronous and optimised for single-thread execution. Both versions are written in the Scala programming language, a relatively new programming language for the Java Virtual Machine. Scala[6] is suited for both object-oriented and functional programming, supports parallel and asynchronous programming and is compatible with Java code and existing libraries.

Both versions are based on the concept of meeting arenas introduced in Section 3.1. Every agent is assigned with a solution to the optimisation problem, its fitness and some "life energy" (a single resource). The behaviour of the agents is the same in both versions (see Listing 1). They differ in how agents join arenas and how arenas execute meetings.

In this evolutionary multi-agent system, we use the following arenas:

- agents are removed from the system in the **death arena**

[6] `http://www.scala-lang.org/`

```
1  def chooseArena = energy match {
2    case 0 => deathArena
3    case e if e > threshold
4           => reproductionArena
5    case e => fightingArena
6  }
```

Listing 1: Agents choose an arena to join based on their current resources, in this case energy.

- agents compare their fitness in the **fighting arena**. Losers give some of their energy to the winners
- new agents are created in the **reproduction arena**. Children solutions are derived from their parents using variation operators. Parents give some of their energy to their children.

5.1. Asynchronous EMAS

This version is similar to the approach in frameworks like Jade, in which agents are the basic unit of concurrency. They are independent entities which do not directly expose state and can only query each other for information.

Agents and arenas have been implemented using the Akka[7] actor library. They are represented by actors which execute asynchronously and communicate through message passing. As such, agents can be mapped to threads in a very flexible way. Akka actors are handled by a component called the *dispatcher*. The dispatcher allows each actor in turn to process one or more messages from its mailbox. It is also used to execute asynchronous tasks.

The processing of a message or task can happen in any thread owned by the dispatcher, which behaviour is fully configurable. The dispatcher can use a single thread, a pool of threads or assign a separate thread to each actor. Akka ensures happens-before relationships between the processing of consecutive messages and preserves memory consistency.

After its previous meeting have ended, every agent chooses an arena and join it by sending a `JoinMeeting` message. Every arena has a fixed size and acts as a cyclic barrier: a meeting is triggered as soon as the capacity of the arena have been reached (see Listing 2). Multiple meetings may be happening at the same time, but every agent can only take part in one of them. When the meeting is finished, a `MeetingEnded` message is sent asynchronously to its participants so that they can choose a new arena to join.

An additional mechanism, omitted above for clarity, triggers a meeting after some inactivity timeout. This may be beneficial for the algorithm (e.g., reproduction with mutation only) or help avoid deadlocks (when the number of agents in the environment is lower than the capacity of the arena).

Listing 3 shows an example of a meeting in the fighting arena. As the arena has no direct access to agents state it needs to query them with the use of

[7] `http://akka.io/`

```
1  def receive = {
2    case JoinMeeting =>
3      waitingRoom.add(sender)
4      if(waitingRoom.isFull()) {
5        val members = waitingRoom.flush()
6        performMeeting(members) andThen {
7          members foreach {
8            member => member tell MeetingEnded
9          }
10       }
11     }
12   }
```

Listing 2: Asynchronous arenas act as a cyclic barrier and trigger an asynchronous meeting as soon as they are full.

messages. A Scala feature known as *futures* and *for comprehension* allows to implement asynchronous and non-blocking meetings. The `askForFitness` and `getEnergyFrom` functions return a *future value* which will be *completed* only when all the members reply to messages. *For comprehension* composes these futures into a new one which is returned from the `performMeeting` function, allowing installation of a *completion hook* (Listing 2, lines 6 to 9). The important thing is that the `performMeeting` function can return before the meeting has actually ended, so that another meeting may be triggered in the arena.

```
1  def performMeeting(members) = for(
2      fitnesses <- askForFitness(members);
3      val winner = zip(members, fitnesses)
4        .maxBy { (m, f) => f }
5        .map { (m, f) => m }
6      val losers = members - winner;
7      energies <- getEnergyFrom(losers)
8      ) yield winner tell ReceiveEnergy(energies.sum)
```

Listing 3: Non-blocking asynchronous fight using Scala futures and *for comprehension*.

5.2. Synchronous EMAS

In this version, agents are considered parts of the model rather than the implementation, like in Netlogo. As such, they are not represented as individual entities but as data structures.

Populations are collections of agents, processed step by step by arenas to yield new collections (see Listing 4). Agents are split into arenas (lines 2-4) and grouped accordingly to the arity of each arena (line 7). Finally, groups are processed by arenas and the agents resulting from each meeting are combined into a new population (lines 5-10).

```
1  def step(population) = {
2    val agentsInArenas = population groupBy { agent =>
3      agent.chooseArena
4    }
5    val newPopulation =  agentsInArenas flatMap {
6      (arena, agents) =>
7        agents grouped(arena.size) flatMap {
8          members => arena.performMeeting(members)
9        }
10   }
11   return newPopulation shuffled
```

Listing 4: Agents are split between arenas, grouped and processed. These action repeatedly transform the population.

The `performMeeting` method of each arena should in this case return a collection of agents representing the result of a meeting. These collections are merged into the new population. The Listing 5 shows the implementations of a synchronous fighting arena, which is similar but simpler than in the asynchronous version, as arenas can now have direct and synchronous access to the state of agents.

```
1  def performMeeting(members) = {
2    val winner = members maxBy {
3      agent => agent.fitness
4    }
5    val losers = members - winner
6    val energies = getEnergyFrom(losers)
7    winner.energy += energies.sum
8    return members
9  }
```

Listing 5: A synchronous fighting arena transforms its members by transferring energy from losers to winners.

The step function from Listing 4 could be executed in a simple loop. However, we used an Akka actor which repeatedly sends a Step message to itself, in order to minimise the performance impact of the Akka framework itself when comparing both versions.

It should be added that the structure of the synchronous version is similar to the MapReduce pattern and could be parallelised in a similar way. While this is a topic of the current research, we decided to stick to a possibly simple version in this work.

6. Experimental Results

We carried out a series of experiments to measure the performance and scalability of the implementations described in the previous section. We applied the evolutionary multi-agent system to the optimisation task of finding the global minimum of the Rastrigin benchmark function (Eq. 1), a highly multimodal function with many local optima and one global minimum equal 0 at $\bar{x} = 0$ (Fig. 3). We used a problem size (a dimension of the function) equal to 100.

$$f(x) = 10n + \sum_{i=1}^{n}(x_i^2 - 10\cos(2\pi x_i)) \tag{1}$$

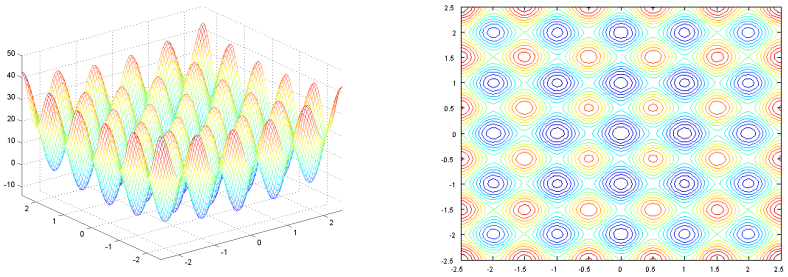

Figure 3: The Rastrigin function in two dimensions.

The parameters used in our experiments are listed in Table 1.

initial-size	50	problem-size	100
initial-energy	10	mutation-rate	0.1
reproduction-threshold	10	mutation-range	0.05
reproduction-transfer	5	mutation-probability	0.75
fight-transfer	10	recombination-probability	0.3
fight-arena-size	2		
migration-probability	0.001		

Table 1: EMAS parameters.

The environment was initialised with *initial-size* agents, each given *initial-energy*. Agents were fighting on an arena of *fight-arena-size* capacity, transferring *fight-transfer* energy from a loser to a winner. As soon as an agent's energy exceeded *reproduction-threshold*, it entered a reproduction arena of size 2. Each pair of agents in the reproduction arena reproduced using a set of genetic operators described below, creating 2 new agents, each one given *reproduction-transfer* energy from one of their parent.

In the second stage of our experiments, at each step every agent had a *migration-probability* of migrating to some other environment. The target environment was chosen at random and including the original one.

In order to create new solutions to be assigned to newborn agents, the following genetic operators were used. Solutions were encoded as real-valued vectors. At each reproduction, crossover and mutation happened with respectively *recombination-probability* and *mutation-probability*. We used random average crossover, which consist in picking a random point in the hypercube defined by the parents genotypes. Every feature in the solution vector was mutated with probability *mutation-rate*. We used gaussian mutation with *mutation-range* standard deviation.

6.1. Performance Testing

In our performance testing, we distinguished four experimental scenarios. All of them share the same set of parameters listed above and have been run on Pl-Grid [8] infrastructure. We used nodes with an Intel Xeon X5650 2,66 GHz processor, with 1GB of memory and a variable number of active cores (up to 12).

The first three scenarios correspond to the asynchronous implementation with an Akka dispatcher configured with respectively *own-thread*, *thread-pool* and *single-thread* policy (see Section 5.1). The fourth scenario is the synchronous implementation (Section 5.2).

Each scenario was repeated 30 times with different random generator seed values. The asynchronous models have been run for 60 minutes each, while the synchronous one only for 10 minutes.

We gathered two metrics: a) the fitness of the best solution found so far at any given time, b) the number of fitness function evaluations at any given time. The former metric shows the efficiency of the evolutionary algorithm itself and is also dependent on i.e. the parameters of the evolutionary operators. The latter reflects the number of agent meetings and only depends on the execution model and threading strategy. Of course, the dynamics of the underlying multi-agent system have an impact on the efficiency of the evolutionary algorithm.

	cores				
	1	2	4	8	12
own	22.1500	19.4800	12.7577	9.2573	3.6087
pool	18.0173	17.6303	11.1574	5.4219	3.8016
single	15.2845	19.2584	6.3323	8.1770	3.4879
sync	0.0371	0.0398	0.0321	0.0186	0.0257

Table 2: Final best fitness found in each of the models with a given number of cores. The asynchronous models were run 60 minutes, the synchronous one was run 10 minutes. The results are averaged over 30 runs.

The results in Tables 2 and 3 indicate that:

[8] `http://www.plgrid.pl/en`

	cores					
	1	2	4	8	12	
own	1.0407	1.5252	1.6836	2.0900	2.9918	$\times 10^7$
pool	1.4270	1.7994	2.0061	2.6123	2.8876	
single	1.7423	1.4949	2.2076	2.2561	2.9611	
sync	3.3296	3.1524	3.7204	5.6854	4.6629	

Table 3: Total number of fitness evaluations in each of the models with a given number of cores. The asynchronous models were run 60 minutes, the synchronous one was run 10 minutes. The results are averaged over 30 runs.

- There was no statistically significant difference in the performance of the asynchronous version using different thread policies (as verified by a two-sample Kolmogorov-Smirnov test with p=0.5).
- The asynchronous versions greatly improved when given more cores...
- ...but were *dramatically* worse than the synchronous version.

This difference in efficiency did not come from a flaw in the evolutionary algorithm itself, but rather from the underlying implementation model. The best asynchronous version only performed about 8×10^3 fitness evaluations per second, while the synchronous version did more than 7×10^4 — nearly an order of magnitude faster. This efficiency gap can clearly be seen in Figures 4 and 5.

Profiling data suggest that the asynchronous implementation was not idle or blocking on I/O, but rather very busy managing threads and passing messages between actors.

Figure 6 shows the empirical distribution functions of the final fitness achieved in separate runs of each model. In many cases, the asynchronous version did converge to acceptable solutions (though given much more time). However, in many runs, they clearly needed more time. In contrast, all the runs of the synchronous version converged to the attraction basin of the global optimum (which corresponds in the case of the Rastrigin function to a fitness value lower than 1). It took an average of 132.13 s ($\pm$ 23.82 s), the empirical cumulative distribution is shown in Figure 7.

6.2. Scalability Testing

Having determined that the synchronous model is more efficient, we went on to test the scalability of the algorithm when new resources were added. We modified the implementation to simultaneously start a synchronous EMAS environment on many nodes in a cluster. The environments discovered each other in the cluster and connected, enabling the migration of agents between them.

We considered several scenarios with an increasing number of nodes. Each scenario was run for 10 minutes and repeated 30 times.

Fig. 8 shows the fitness of best solution found after a given time, averaged over all environments in a given scenario and all runs of the scenario. We can see that even adding a second node to the computation leads to significantly better results. Moreover, adding more nodes increases the convergence rate.

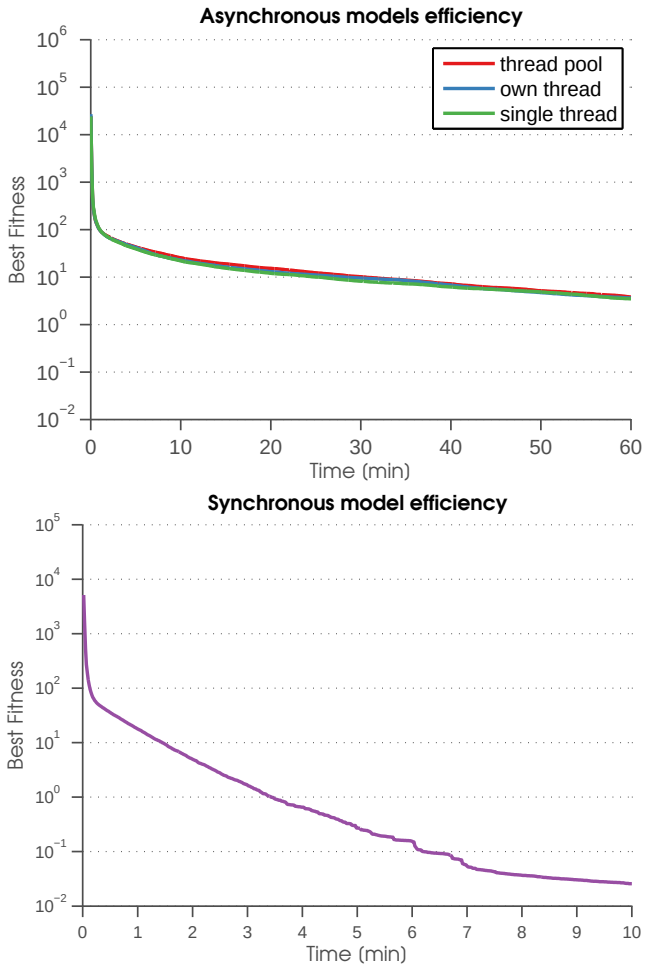

Figure 4: Average best fitness found in each of the models after a given amount of time (for 12 cores used). Top – asynchronous models, bottom – synchronous model.

These results show that is efficient to decompose multi-agent systems into distributed environments, like in the classical island model of evolutionary algorithms. However, the decentralised semantics of agent interaction may lead to more intelligent migration strategies, for example where agent populations automatically balance the load in the cluster.

7. Conclusion

Metaheuristics can be valuable in decision support systems with time constraints. We discussed in this paper how the agent approach can be applied to these systems in order to build efficient and scalable software. We described

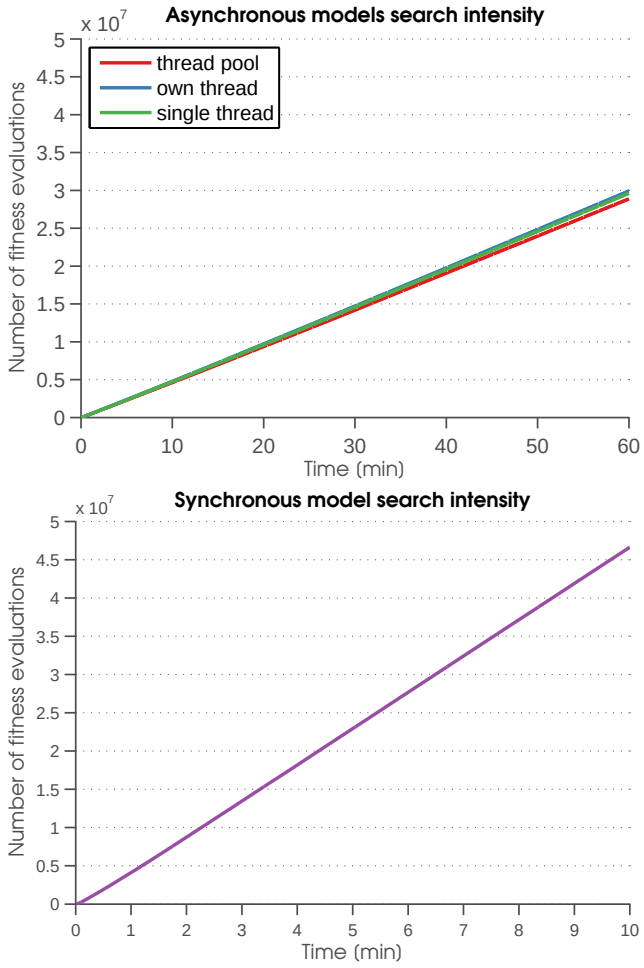

Figure 5: Amount of fitness evaluations performed in each of the models after a given amount of time (for 12 cores used). Top – asynchronous models, bottom – synchronous model.

the concept of evolutionary multi-agent systems, an example of a metaheuristic combining agent-based and evolutionary techniques.

The main goal of our work was to investigate the existing methods of building agent software and suggest new directions of development. In particular, we wanted to see if the dominant approach, which considers every agent as a unit of concurrency, is really efficient in computational intensive simulations.

To that purpose, we developed two alternate implementations of an evolutionary multi-agent system which generalise the trends in existing agent software. We introduced the idea of meeting arenas, an agent-based realisation of the Mediator design pattern which allow to efficiently structure multi-agent systems. We applied this concept to two versions of the algorithm: an asyn-

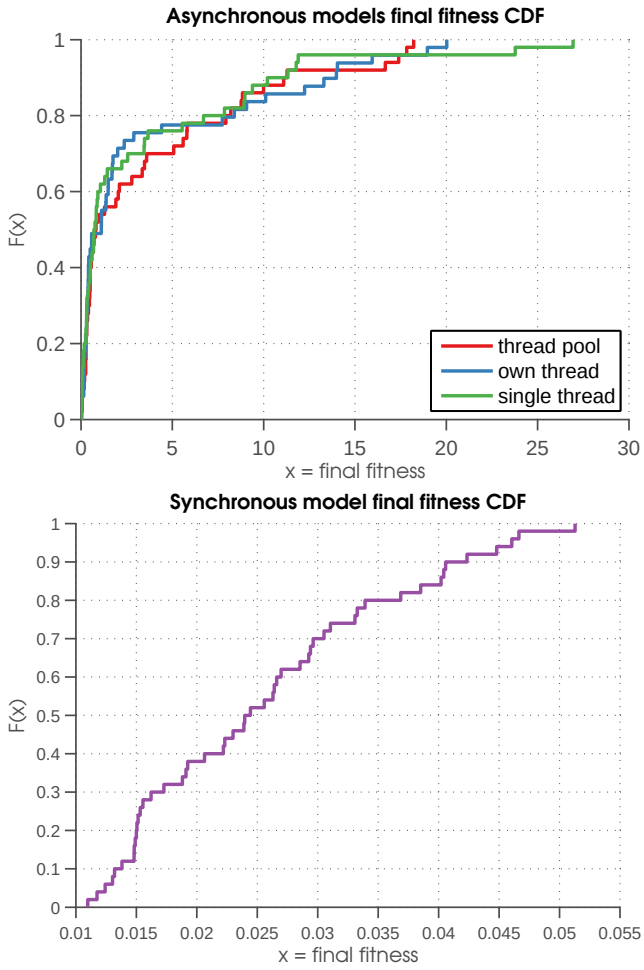

Figure 6: Empirical cumulative distribution functions of the final fitness values found at the end of the computations (for 12 cores used). Top – asynchronous models, bottom – synchronous model.

chronous one, where every agent is a fully independent entity, and a synchronous one which treats agents as simple data structures.

Our experiments revealed that an asynchronous implementation, which may feel more *the agent way*, is nearly an order of magnitude less efficient than a synchronous one based on the same design. Several prototypes in other technologies supported these results. Further experiments on the synchronous implementation demonstrated that it can easily be scaled in a distributed setting, so that the efficiency of the algorithm increases when new nodes are added to the computation.

Therefore, we showed that the prevailing approach in existing agent plat-

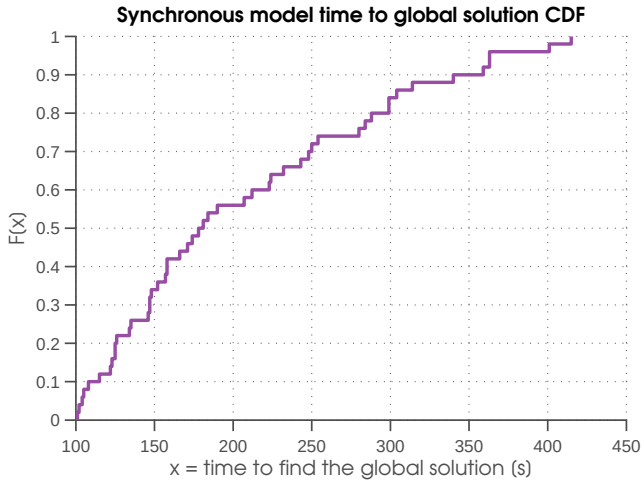

Figure 7: Empirical cumulative distribution function of the time required to find the global solution in the synchronous model (for 12 cores used).

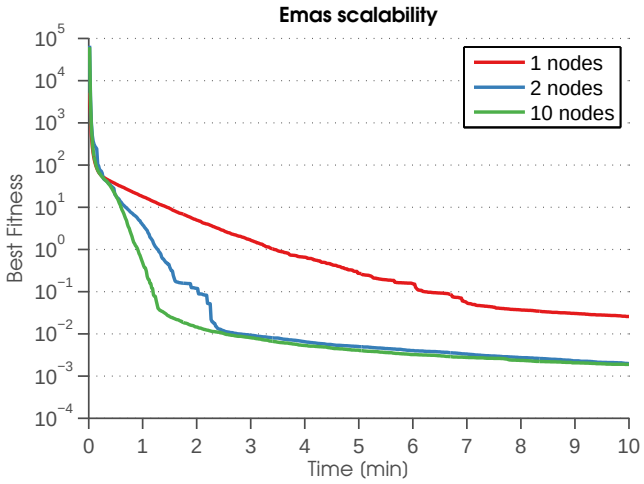

Figure 8: Best fitness reached within a given time when using different amounts of nodes (parallel environments).

forms is not best suited in this particular class of applications. Instead, there is still room for improvement in the field of agent software dedicated to intensive simulations and computations. To that purpose, the concept of meeting arenas introduced in this paper allow to retain the expressive power of existing agent-based algorithms but can lead to much more efficient synchronous implementations.

In the nearest future, we want to see if concepts used in the functional

programming paradigm could be more suited or more efficient in multi-agent software than the dominant object-oriented approach. Future work could also tell what kind of parallelism could be efficiently introduced in populations of agents. In particular, it would also be interesting to see how the Map/Reduce paradigm could be used to develop efficient massive multi-agent systems with hundreds of thousands of agents.

References

[1] Z. Michalewicz, D. Fogel, How to Solve It: Modern Heuristics, Springer, 2004.

[2] R. Dreżewski, J. Sepielak, L. Siwik, Classical and agent-based evolutionary algorithms for investment strategies generation, in: A. Brabazon, M. O'Neill (Eds.), Natural Computing in Computational Finance, volume 185 of *Studies in Computational Intelligence*, Springer-Verlag, 2009, pp. 181–205.

[3] J. Koźlak, G. Dobrowolski, M. Kisiel-Dorohinicki, E. Nawarecki, Anti-crisis management of city traffic using agent-based approach, Journal of Universal Computer Science 14 (2008).

[4] A. Ławrynowicz, A survey of evolutionary algorithms for production and logistics optimization, Research in Logistics & Production 1 (2011).

[5] M. Wooldridge, N. Jennings, Intelligent agents, in: LNAI 890, Springer Verlag, 1995.

[6] Y. Luo, K. Liu, D. Davis, A multi-agent decision support system for stock trading, IEEE Network Jan/Feb (2002) 20–27.

[7] S. Ossowski, A. Fernandez, J. M. Serrano, J. Perez-de-la Cruz, M. Belmonte, J. Hernandez, A. Garcia-Serrano, J. Maseda, Designing multiagent decision support system-the case of transportation management, in: Proc. of AAMAS, ACM Press, 2004, pp. 1469–1469.

[8] F. Bellifemine, A. Poggi, G. Rimassa, JADE: a FIPA2000 compliant agent development environment, in: Proceedings of the fifth international conference on Autonomous agents, ACM, 2001, pp. 216–217.

[9] M. J. North, N. T. Collier, J. Ozik, E. R. Tatara, C. M. Macal, M. Bragen, P. Sydelko, Complex adaptive systems modeling with Repast Simphony, Complex Adaptive Systems Modeling 1 (2013) 3.

[10] O. Gutknecht, J. Ferber, The madkit agent platform architecture, in: T. Wagner, O. Rana (Eds.), Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems, 2001.

[11] K. Pietak, A. Woś, A. Byrski, M. Kisiel-Dorohinicki, Functional integrity of multi-agent computational system supported by component-based implementation, in: Proceedings of the 4th International Conference on Industrial Applications of Holonic and Multi-agent Systems, 2009.

[12] A. Byrski, R. Debski, M. Kisiel-Dorohinicki, Agent-based computing in an augmented cloud environment, Computer Systems Science and Engineering 27 (2012).

[13] Ł. Faber, K. Pietak, A. Byrski, M. Kisiel-Dorohinicki, Agent-Based Simulation in AgE Framework, in: A. Byrski, Z. Oplatkova, M. Carvalho, M. Kisiel-Dorohinicki (Eds.), Advances in Intelligent Modelling and Simulation, volume 416 of *Studies in Computational Intelligence*, Springer Berlin Heidelberg, 2012, pp. 55–83.

[14] A. Byrski, R. Dreżewski, L. Siwik, M. Kisiel-Dorohinicki, Evolutionary multi-agent systems, The Knowledge Engineering Review 30 (2015 (in press)).

[15] D. Power, Decision support systems: concepts and resources for managers, Quorum Books, 2002.

[16] Z. Michalewicz, Ubiquity symposium: Evolutionary computation and the processes of life: the emperor is naked: evolutionary algorithms for real-world applications, Ubiquity 2012 (2012) 3:1–3:13.

[17] F. Glover, K. G.A., Handbook of Metaheuristics, Springer, 2003.

[18] J. Dréo, A. Pétrowski, P. Siarry, E. Taillard, A. Chatterjee, Metaheuristics for Hard Optimization: Methods and Case Studies, Springer, 2005.

[19] E.-G. Talbi, Metaheuristics: From Design to Implementation, Wiley, 2009.

[20] M. Kisiel-Dorohinicki, G. Dobrowolski, E. Nawarecki, Agent populations as computational intelligence, in: L. Rutkowski, J. Kacprzyk (Eds.), Neural Networks and Soft Computing, Physica Verlag, 2002, pp. 608–614.

[21] K. Cetnarowicz, M. Kisiel-Dorohinicki, E. Nawarecki, The application of evolution process in multi-agent world (MAW) to the prediction system, in: M. Tokoro (Ed.), Proc. of the 2nd Int. Conf. on Multi-Agent Systems (ICMAS'96), AAAI Press, 1996.

[22] A. Byrski, Tuning of agent-based computing, Computer Science 14 (2013).

[23] A. Byrski, M. Kisiel-Dorohinicki, Immunological selection mechanism in agent-based evolutionary computation, in: Proc. of IIS: IIPWM '05 conference : Gdansk, Poland, Advances in Soft Computing, Springer Verlag, 2005, pp. 411–415.

[24] D. Krzywicki, Niching in evolutionary multi-agent systems, Computer Science 14 (2013).

[25] G. Danoy, P. Bouvry, O. Boissier, A multi-agent organizational framework for coevolutionary optimization, in: Transactions on Petri nets and other models of concurrency IV, Springer, 2010, pp. 199–224.

[26] M. Milano, A. Roli, Magma: a multiagent architecture for metaheuristics., IEEE transactions on systems, man, and cybernetics. Part B, Cybernetics 34 (2004) 925–941.

[27] E. Noda, A. L. Coelho, I. L. Ricarte, A. Yamakami, A. A. Freitas, Devising adaptive migration policies for cooperative distributed genetic algorithms, in: Systems, Man and Cybernetics, 2002 IEEE International Conference on, volume 6, IEEE, 2002, pp. 6–pp.

[28] D. Meignan, J.-C. Créput, A. Koukam, An organizational view of metaheuristics, in: First International Workshop on Optimisation in Multi-Agent Systems, AAMAS, volume 8, 2008, pp. 77–85.

[29] L. Braubach, A. Pokahr, Developing Distributed Systems with Active Components and Jadex, Scalable Computing: Practice and Experience 13 (2012) 100–119.

[30] C. Frantz, M. Nowostawski, M. Purvis, Dynamic ad hoc coordination of distributed tasks using micro-agents, in: Agents in Principle, Agents in Practice, 2011, pp. 275–286.

[31] M. Ughetti, T. Trucco, D. Gotta, Development of Agent-Based, Peer-to-Peer Mobile Applications on ANDROID with JADE, 2008 The Second International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (2008) 287–294.

ABOUT THE AUTHORS

Daniel Krzywicki obtained his M.Sc. in 2012 at AGH University of Science and Technology in Cracow and is currently a Ph.D. student at the Department of Computer Science of AGH-UST. His research interests include agent-based computations, functional programming and distributed systems.

Łukasz Faber obtained his M.Sc. in 2012 at AGH University of Science and Technology in Cracow and is currently a Ph.D. student at the Department of Computer Science of AGH-UST. His research interests include agent-based modelling and distributed systems.

Aleksander Byrski obtained his Ph.D. in 2007 and D.Sc. (habilitation) in 2013 at AGH University of Science and Technology in Cracow. He works as an assistant professor at the Department of Computer Science of AGH-UST. His research focuses on multi-agent systems, biologically-inspired computing and other soft computing methods.

Marek Kisiel-Dorohinicki obtained his Ph.D. in 2001 and D.Sc. (habilitation) in 2013 at AGH University of Science and Technology in Cracow. He works as

an assistant professor at the Department of Computer Science of AGH-UST. His research focuses on intelligent software systems, particularly using agent technology and evolutionary algorithms, but also other soft computing techniques.

ACKNOWLEDGEMENT

The research presented in the paper was partially supported by the European Commission FP7 through the project ParaPhrase: Parallel Patterns for Adaptive Heterogeneous Multicore Systems, under contract no.: 288570 (http://paraphrase-ict.eu).

The research presented in the paper was conducted using PL-Grid Infrastructure (http://www.plgrid.pl/en).